Adv. Math. Econ. 8, 135–153 (2006)

Advances in
MATHEMATICAL
ECONOMICS

A game-theoretic approach to global warming

Prajit K. Dutta[1] **and Roy Radner**[2]

[1] Dept. of Economics, Columbia University
(e-mail: pkd1@columbia.edu)
[2] Stern School of Business, New York University
(e-mail: rrandner@stern.nyu.edu)

Received: August 21, 2005
Revised: October 14, 2005

JEL classification: C73, Q54, Q58

Mathematics Subject Classification (2000): 91A25, 91A80

Abstract. In the absence of a world government, stopping the advance of global warming requires implementation of self-enforcing treaties among the countries of the world. In the language of game theory, such treaties are Nash equilibria of an underlying dynamic "climate change game." In this paper, we report on the progress of a project to formulate and analyze models of such a game. The players are the sovereign countries of the world (say the roughly 200 members of the United Nations). The rules of this game are determined by the laws of physics and chemistry, and by the economic resources of the various countries. An important property of our models is the large multiplicity of equilibria. Indeed, this property enables us to find "Pareto-improving" equilibria, i.e., that improve the outcome for every country relative to the "business-as-usual equilibrium" we seem to be in at the present time. In each model we describe the set of equilibria, the business-as-usual equilibrium, and equilibria that are Pareto-improving relative to business-as-usual. Since much of the global warming is caused by the accumulation of greenhouse gases (GHGs) in the earth's atmosphere, and the GHGs dissipate very slowly, an appropriate model must be in the form of a dynamic game, with state variables that change over time as a consequence of the actions of the individual countries. Thus, the state variables include the global stock of GHG and the state of the relevant technology in each country.

1. Introduction

It is now generally recognized, at least in the scientific community, that global warming presents a significant threat to the environment of the earth, and that human activity in the past two centuries is a significant factor in this phenomenon, largely through the increased burning of fossil fuels. (For background material see Dutta and Radner, 2004, 2005a, the references cited there, and Raven, 2005.) The term "global warming" is, of course, an oversimplification of a complex process of climate change comprising trends of increased average temperature, increased climate and weather volatility, and local changes in climate patterns involving both warming and cooling. This process, if left unchecked, threatens to impose large, if still uncertain, costs on various parts of the world, especially those with poorer populations.

In the absence of a world government, checking the advance of this process requires the implementation of self-enforcing treatiesamong the countries of the world. In the language of game theory, such treaties are Nash equilibria of an underlying vast dynamic "climate change game."

In this paper, we report on the progress of a project to formulate and analyze models of such a game. In these models, the players are taken to be the sovereign countries of the world (say the roughly 200 members of the United Nations). The rules of this game are determined by the laws of physics and chemistry, and by the economic resources of the various countries. (See, e.g., Radner, 1999, for a discussion of the contrast of the situation represented by the climate-change game and the typical situation modeled in the "mechanism design" literature.) It is, of course, a heroic assumption to hypothesize that the individual countries of the world are capable of the rational behavior that is assumed in the theory of games. Indeed, the behavior of each sovereign country can itself be viewed as the outcome of a political game played by the country's inhabitants. In defense of our approach we can only plead that taking the players of the game to be the individual inhabitants of the world would make the theoretical analysis intractable, at least at this stage of our understanding. Thus we regard our approach as part of a reasonable decomposition of the global problem into a set of more tractable subproblems.

An important property of our game-theoretic models of the climate-change game is the large multiplicity of equilibria. Indeed, this property gives us hope that we can find equilibria (self-enforcing treaties) that improve the lot of every (or most) countries relative to the "bad equilibrium" that we seem to be at the present time. In the language of game theory, we are searching for "Pareto-improving" equilibria. In this project, we have formulated several models of increasing complexity, and in each case have tried to describe the set of equilibria, the one we seem to be in now (which we call "business-as-usual"), and at least some of the equilibria that are Pareto-improving relative to business-as-usual.

In Section 2 we formulate the basic model. Since much of the global warming is caused by the accumulation of greenhouse gases (GHGs) in the earth's atmosphere, and the GHGs dissipate very slowly, an appropriate model of the climate change game must be in the form of a dynamic game, with state variables that change over time as a consequence of the actions of the individual countries. In particular, one of the actions of each country in each period is the quantity of GHGs emitted into the atmosphere. Most of the GHG is in the form of carbon-dioxide formed in processes of burning fossil fuels to produce energy, which in turn is an input into the production of other goods and services. Countries can reduce their emissions of GHGs by reducing their total gross domestic product (GDP) per period, and thus their consumption of energy, and/or by modifying their production technology to economize on enegy and — even more importantly — to reduce the amount of GHG emitted per unit of energy produced and consumed. Thus, the state variables include the global stock of GHG and the state of the relevant technology in each country. (We use the stock of GHG as a surrogate for the degree of global warming.) We define the payoff to each country to be the total discounted GDP less the sum of the cost of improving the technology and the cost of the damage done by the stock of GHG.

For any given initial values of the state variable, there is a set of equilibria of the ensuing dynamic game, which defines the *equilibrium correspondence*, and a set of Pareto-optimal outcomes (trajectories), parametrized by the weights assigned to each country in the global welfare function, which defines the Pareto-optimum correspondence. In Section 3 we provide information about these two correspondences. In particular we characterize the business-as-usual (BAU) equilibria, and show that they cannot be Pareto-optimal. We also provide some conditions under which there are equilibria that are Pareto-superior to the business-as-usual equilibrium, i.e., better than the BAU equilibrium for every country. These are, of course, the candidates for self-enforcing treaties.

In Section 4 we consider the special case in which the emission-producing technology is fixed, thus leaving the stock of GHG as the only state variable. For this case we can give a complete characterization of the equilibrium correspondence, together with a discussion of some interesting particular equilibria. Section 5 provides some brief concluding remarks about other models we have studied, and some future research. Section 6 contains the references cited in the paper. There we have relied on a small set of key references that point the reader to relevant literature.

In preparing this summary we have borrowed heavily from other papers produced in the course of this project (see Section 6 for more detail). Also, although our formulation of the models and statements of the theorems are relatively self-contained, the reader is referred to the original papers for all proofs.

2. A climate-change game: the basic model

In this section we describe a mathematical model of a *dynamic climate-change game*. The players in the game are countries, and it is assumed that each country has the authority and political will to control its own rate of emission of greenhouse gases, subject to technological and resource constraints.

There are I countries. The emission of (a scalar index of) greenhouse gases (GHGs) during period t by country i is denoted by $a_i(t)$. [Time is discrete, with $t = 0, 1, 2, \ldots$, ad inf., and the $a_i(t)$ are nonnegative.] The emission of GHG in each country is related to its level of economic activity, notably the production and use of energy produced by burning fossil fuels, although there are other sources of GHGs. For simplicity we let $e_i(t)$ denote a scalar index of inputs into production and consumption associated with the emission of GHGs during period t by country i. For brevity, we shall call $e_i(t)$ the level of "energy input." The output of the country is described by a scalar index, e.g., "gross domestic product" (GDP). This output depends on $e_i(t)$ and other inputs according to the country's current "production function." In our model, the production function is in a "reduced form," implicitly reflecting for each level of $e_i(t)$ the corresponding levels of the other inputs and can be interpreted as holding constant in time, for each country, its stocks of capital and labor, and the technology of production, except for the production of energy. Thus country i's GDP in period t is denoted by $Y_i[e_i(t)]$. Given the country's current technology, it's emission of GHG during the period is assumed to be

$$a_i(t) = f_i(t)e_i(t), \qquad i = 1, \ldots, I. \tag{1}$$

The coefficient $f_i(t)$ will be called the *emission factor* of country i in period t. [In an equivalent model, the emission factors are constant in time, but every unit of GDP may be produced (more efficiently) by successively smaller amounts of energy as time passes.]

Let $A(t)$ denote the global (total) emission during period t;

$$A(t) = \sum_{i=1}^{I} a_i(t). \tag{2}$$

The total (global) stock of GHGs at the beginning of period t is denoted by $g(t) + g_0$, where g_0 is what the "normal" steady-state stock of GHGs would be if there were negligible emissions from human sources (e.g., the level of GHGs in the year 1800). We might call $g(t)$ the *excess GHG*, but we shall usually suppress the word "excess." The law of motion for the GHG is assumed to follow the linear difference equation,

$$g(t+1) = A(t) + \sigma g(t), \tag{3}$$

where σ is a given parameter $(0 < \sigma < 1)$. (This linear approximation is a gross simplification of greenhouse gas dynamics, but reasonable for the present game-theoretic problem.) We may interpret $(1 - \sigma)$ as the fraction of the beginning-of-period stock of GHG that is dissipated from the atmosphere during the period. The "surviving" stock, $\sigma g(t)$, is augmented by the quantity of global emissions, $A(t)$, during the same period.

We assume that for each country the cost of the damage due to climate change is linear in the global stock of GHG, i.e., equal to $c_i g(t)$, and is subtracted from the country's gross domestic product in that period.

Finally, each country can reduce its own emission factor, $f_i(t)$, but at a cost. We assume that this cost is proportional to the decrease in the emission factor, i.e., equal to $k_i[f_i(t) - f_i(t+1)]$. Actions taken in one period to reduce its emission factor take effect in the next period. We assume that $k_i > 0$, and that the changes in the emission factors are constrained by

$$m_i \leq f_i(t+1) \leq f_i(t). \tag{4}$$

Thus in each period a country can only reduce its emission factor, not increase it, and there is a lower bound on the (eventual) level of its emission factor.

The *utility* of country i in period t (*one-period payoff*) is

$$v_i(t) = Y_i[e_i(t)] - c_i g(t) - k_i[f_i(t) - f_i(t+1)]. \tag{5}$$

Let δ denote the discount factor; then the total discounted utility (*total payoff*) for country i is

$$v_i = (1-\delta)\sum_{t=0}^{\infty} \delta^t v_i(t), \qquad i = 1, \ldots, I. \tag{6}$$

Note that each country's current-period payoff depends directly on its *current* energy usage only through its production function, but also depends on its own and others' *previous* energy usage and emission factors through their effects on the current stock of GHG. Note, also, that the present value in the payoff function has been *normalized*, so that as the discount factor approaches unity the normalized present value will typically approach as a limit the *long-run average payoff*. This does not impact the analysis when the discount factor is fixed, but it does influence the interpretation of the numerical results when the discount factor is varied (Section 4).

We assume that Y_i is strictly concave and twice differentiable, and reaches a maximum at some finite level of energy use. The damage cost coefficients, c_i, are constant in time and strictly positive ($c_i > 0$), although our method of analysis would allow them to have either sign. The discount factor, δ, is the same for all countries, with $0 < \delta < 1$.

The state of the system at the beginning of period t is characterized by the $(I+1)$-dimensional vector, $s(t) = [f(t), g(t)]$, where $f(t) = [f_1(t), \ldots, f_I(t)]$. A *strategy* for a country determines for each period the country's energy usage and emission factor as a function of the entire past history of the system, including the state variables up to the current period and the past actions of all the countries. A *Nash Equilibrium* is a profile of strategies such that no individual country can increase its payoff by *unilaterally* changing its strategy. A Nash equilibrium is the formal construct that corresponds to a self-enforcing treaty. There will typically be many Nash equilibria of the climate change game, and the set of Nash equilibria will depend on the initial state of the system. It is important to note that *there is no way for a country in any period to commit itself to follow a particular strategy in the future*. In particular, since there is no world government, countries cannot sign binding contracts.

A *stationary strategy* for country i is a strategy that is history-independent and only depends on the current state, s, which is then mapped into a current action, a_i. A *Markov Nash Equilibrium (MNE)* is a Nash Equilibrium in which every country's strategy is stationary. In a Markov Nash Equilibrium, *no matter which period and history of emissions we consider, a country's best option from that point on is to follow through on the remainder of its Markov Nash strategy*.

Finally, it is useful to have as a benchmark the concept of a *global Pareto optimum*. Let $x = (x_i)$ be a vector of positive numbers, one for each country. A *global Pareto optimum (GPO)* corresponding to x is a profile of strategies that maximizes the weighted sum of country payoffs,

$$v = \sum_i x_i v_i, \tag{7}$$

which we shall call the *global welfare*. One interpretation of a GPO is that it is what a "world government" would like to do for the world if it could force the national governments to act in the way that it deemed fit.

Without loss of generality, we may take the weights, x_i, to sum to I. We emphasize that to each vector of weights there corresponds a different global welfare function, and hence (in general) a different GPO.

3. Equilibrium and optimal time-paths

3.1 Business-As-Usual equilibrium

In this section we provide information about the set of equilibrium trajectories and the set of global Pareto optimal (GPO) trajectories. We start by describing a particular equilibrium, which we call *Business as Usual* (BAU). This benchmark equilibrium appears to correspond to what we currently observe in the world.

We shall show that BAU equilibrium strategies have the form:

$$e_i(t) = E_i[f_i(t)], \quad f_i(t+1) = F_i[f_i(t)], \qquad t \geq 0,\ i = 1, \ldots, I. \tag{8}$$

Furthermore, $f_i(t)$ will be constant after period 0, and will equal either m_i or $f_i(0)$. Note that such a strategy is stationary, as defined in the previous section. In fact, *the argument of a BAU strategy is the country's own current emission factor only, and does not include the current stock of greenhouse gas or the emission factors of the other countries.*

Here is a precise characterization of BAU strategies. Define

$$w_i = \frac{c_i}{1 - \delta\sigma}. \tag{9}$$

Assume that, for each country,

$$Y_i'(0) > \delta w_i f_i(0). \tag{10}$$

Define the function E_i implicitly by the equation,

$$Y_i'[E_i(y)] = \delta w_i y. \tag{11}$$

Define the function Z_i by:

$$Z_i(y) = k_i y + \left(\frac{\delta}{1-\delta}\right)\{Y_i[E_i(y)] - \delta w_i y E_i(y)\}, \tag{12}$$

and let $F_i(f_i)$ be a value of y that maximizes $Z_i(y)$ subject to the constraint corresponding to (4), i.e.,

$$F_i(f_i) = \arg\max_y \{Z_i(y) : m_i \leq y \leq f_i\}. \tag{13}$$

(If there is more than one maximizing value of y, pick any one.) Note that, although the function Z_i does not depend on f_i, the function F_i does, because of the constraint. We shall call

$$(E, F) = \{(E_i, F_i) : i = 1, \ldots, I\}$$

a *BAU strategy-profile.*

Observe that, since Y_i is concave, *$E_i(y)$ is decreasing in y* (use (11)). Using (12), one verifies that

$$Z_i'(y) = k - \left(\frac{\delta^2 w_i}{1-\delta}\right) E_i(y),$$

and hence $Z_i(y)$ is a *convex function of y*. It follows that $Z_i(y)$ is maximized in y at one of the end-points of the interval $[m_i, f_i]$. There are two cases to consider. If $Z_i'(m_i) \geq 0$, then $F_i(f_i) = f_i$ for all $f_i \geq m_i$. If $Z_i'(m_i) < 0$, then there is some $y_i^0 > m_i$ such that

$$F_i(f_i) = m_i \quad \text{for } f_i < y_i^0,$$
$$F_i(f_i) = f_i \quad \text{for } f_i > y_i^0.$$

(Thus F_i is a "bang-bang" policy.) Note that *each emission factor is constant after period* 1.

Let $V_i(f, g)$ denote country i's (total discounted) payoff when each country uses its BAU strategy and the initial state is $(f, g) = [f(0), g(0)]$. The function V_i is called country i's *value function.*

Theorem 1. *A BAU strategy profile is a Markov equilibrium, called a BAU equilibrium. Along the equilibrium path, each country i's emission and emission factor are constant after the first period, and the emission factor is equal either to m_i or $f_i(0)$. The value function for country i is*

$$V_i(f, g) = Y_i[E_i(f_i)] - c_i g - k_i[f_i - F_i(f_i)] + u_i - \delta w_i g', \tag{14}$$

where

$$u_i = \left(\frac{\delta}{1-\delta}\right)\left\{Y_i[E_i(F_i(f_i))] - \delta w_i \sum_j F_j(f_j)E_j[F_j(f_j)]\right\}, \tag{15}$$
$$g' = \sigma g + \sum_j f_j E_j(f_j).$$

The proof of this theorem uses a standard dynamic programming method, and can be found in (Radner, 1999).

3.2 Global Pareto optimal strategy profiles

We now characterize the global Pareto optimal (GPO) strategy profiles, for the same underlying model. As in (7), let $x_i > 0$ be the weight given to country i in the global welfare function. Define the emission policy function, $\hat{E}_i$, implicitly by the equation,

$$Y_i'[\hat{E}_i(y)] = \frac{\delta w y}{x_i}, \tag{16}$$

where

$$w = \frac{\sum_i x_i c_i}{1 - \delta\sigma} = \sum_i x_i w_i. \tag{17}$$

Define the functions $\hat{Z}_i$ by

$$\hat{Z}_i(y) = k_i y + \left(\frac{\delta}{1-\delta}\right)\left\{Y_i[\hat{E}_i(y)] - \delta\left(\frac{w}{x_i}\right) y \hat{E}_i(y)\right\}, \tag{18}$$

and let $\hat{F}_i(f_i)$ maximize $\hat{Z}_i(y)$ subject to the constraint

$$m_i \leq y \leq f_i. \tag{19}$$

Comparing (18) with (12) reveals that the coefficient w_i in (12) has been replaced by (w/x_i). Again, one can show that $\hat{Z}_i(y)$ is convex in y, and hence $\hat{F}_i(f_i)$ equals either m_i or f_i, according as $\hat{Z}_i(m_i)$ is $>$ or $<$ $\hat{Z}_i(f_i)$.

Theorem 2. *Given the weights (x_i), the following strategy profile is globally Pareto optimal (GPO): for each country i,*

$$\begin{aligned} e_i(t) &= \hat{E}_i[f_i(t)], \\ f_i(t+1) &= \hat{F}_i[f_i(t)]. \end{aligned} \tag{20}$$

In particular, both the emission factor and the "energy input" are constant after the initial date.

The proof is similar to that of Theorem 1.

3.3 A comparison of BAU and GPO strategy profiles

As one might expect, a BAU equilibrium need not be optimal, not only because "energy" inputs are too high, but also because a country that should reduce its emission factor in a Pareto optimum may not do so in the BAU. However, the comparison of the BAU with the set of GPOs is not as straightforward as one might at first expect.

We start with a useful lemma, which implies that if, for a given country i in a given period, the country uses the same emission factor in the BAU and GPO profiles, then its BAU energy use (and hence emission) exceeds its GPO energy use (and hence emission).

Lemma 3. *Let E_i and $\hat{E}_i$ be the BAU and GPO emission functions defined in (11) and (16), respectively. For all vectors $x = (x_i)$ of strictly positive weights, and every emission factor $f_i > 0$,*

$$E_i(f_i) > \hat{E}_i(f_i). \tag{21}$$

An important implication of the preceding lemma is:

Theorem 4. *A Business-As-Usual equilibrium is not globally Pareto optimal.*

[For proofs of the preceding lemma and theorem, see (Dutta and Radner, 2004).]

Does moving from a BAU to a corresponding *non-constrained* GPO necessarily reduce the *emissions* of all countries? The answer to this question is more complex. Recall that we have assumed that emission factors can be decreased, but not increased. (See the final section for remarks on this assumption.) Somewhat paradoxically, *a decrease in a country's emission factor need not lead to a decrease in GPO emissions for that country*. Hence, a decrease in one country's emission factor could lead to a loss in welfare for other countries. To see this, consider a GPO profile. If at a given point of time the emission factor of a country is f_i, then the GPO emission is

$$\hat{a}_i = f_i \hat{E}_i(f_i),$$

so that

$$\frac{d\hat{a}_i}{df_i} = \hat{E}_i(f_i) + f_i \hat{E}_i'(f_i).$$

Of course, if the energy input were held constant, then a decrease in the emission factor would result in a decrease in the emissions. However, we have already seen that

$$\hat{E}_i'(f_i) < 0. \tag{22}$$

Hence a decrease in the emission factor for a given country has two opposing effects. From the preceding inequality,

$$\frac{d\hat{a}_i}{df_i} > 0$$

if and only if

$$\left(\frac{d\log \hat{E}_i(f_i)}{d\log f_i}\right) = \hat{E}_i'(f_i)\left[\frac{f_i}{\hat{E}_i(f_i)}\right] > -1. \tag{23}$$

Note that the absolute value of the left-hand-side of (23) is what economists call the *elasticity* of $\hat{E}_i(f_i)$ with respect to f_i.

The next theorem states that, if the preceding inequality is satisfied, then the switch from a BAU to a corresponding GPO will not increase any country's emission factor, but may decrease it.

Theorem 5. *Suppose that* (23) *is satisfied; then*

$$\textit{if}\ \ F_i(f_i) = m_i, \ \ \textit{then}\ \ \hat{F}(f_i) = m_i, \ \ \textit{whereas}$$
$$\textit{if}\ \ F_i(f_i) = f_i, \ \ \textit{then}\ \ \hat{F}(f_i) = m_i \ \ \textit{or}\ \ f_i.$$

Thus, the switch from the BAU to a GPO with the same initial state will not increase any country's emission factor in any period, but may decrease it.

[For a proof of the theorem, see (Dutta and Radner, 2002).]

As an immediate corollary of the preceding lemma and theorem, we have:

Corollary 6. *If* (23) *is satisfied, then the switch from the BAU to a GPO with the same initial state will decrease every country's emissions in every period.*

3.4 The set of equilibrium outcomes — two results

The previous subsection described one particular equilibrium, "Business as Usual," which we think of as characterizing the current world situation. There are, however, infinitely many equilibria, to which we now turn our attention. Although we do not have a complete characterization of the set of all equilibria, we can give some useful information about them. In particular, we can show that there are cases (sets of parameter values) for which there exist equilibria that are Pareto-superior to the BAU, i.e., better for every country, provided that the countries' discount factor is not too low. On the other hand, there are cases for which no global Pareto optimum can be sustained by an equilibrium. Our treatment here is informal. [See (Dutta and Radner, 2002, 2005a) for a fuller treatment.] The basic message is that, even if GPOs cannot be sustained by equilibria, typically there will be equilibria that are Pareto superior to the BAU, and that set of equilibria will be larger, the closer is the discount factor to unity.

We start with a negative result. To sustain a GPO, the countries must credibly be able to threaten to increase their emissions above the GPO level in the event of a "defection." Suppose that in (4), $m_i = 0$ for all i, and that in a GPO every country reduces its emission factor to zero. If all the countries but one follow the GPO, then it will not be possible for them to punish a defector, because their emission factors will already be zero. Hence it will not be possible to sustain the GPO as an equilibrium. This result is still valid with a small amount of convexity in the technical-change cost function.

Here is a sketch of a positive result. Suppose that the minimum attainable emission factors are all strictly positive. (In (4), $m_i > 0$ for all i.) For any vector of weights, there is a discount factor sufficiently close to unity such that the corresponding GPO requires that every country reduce its emission factor to the minimum. For some set of weights, the GPO is Pareto-superior to the BAU when the emission factors are at their minimum. Hence, for δ sufficiently close to 1, the GPO can be sustained by a "trigger strategy" in which the players threaten to revert to the BAU (given the then current state) in the case of a defection.

4. The case of fixed technology

In this section we consider the special case of the "basic model" of Section 2 in which the emissions-producing technology is fixed, i.e., the emission factor for each country, f_i, is fixed and constant through time. Thus the only strategic variable for country i in each time period t is the level of energy input, or equivalently, the level of emission, $a_i(t)$, as one sees from equation (1). To simplify the notation, corresponding to equation (5), denote the utility of country i in period t as

$$v_i(t) = h_i[a_i(t)] - c_i g(t), \tag{24}$$

where

$$h_i[a_i(t)] = Y_i[e_i(t)]$$

is the GDP of country i in period t if it's emission level in that period is $e_i(t)$. Note that, compared to equation (5), there is no term corresponding to the cost of changing the emission factor.

For this special case, we can obtain much more information about the set of subgame-perfect equilibria (SPEs). In particular, in what follows we shall (1) characterize the equilibrium payoff correspondence, (2) describe "extreme equilibria," including the second best SPEs, and the "worst" SPE for each country, and (3) describe a family of SPEs that we call "greenhouse trap equilibria."

4.1 The equilibrium payoff correspondence

We can show that the SPE payoff correspondence has a surprising simplicity; the *set* of equilibrium payoffs at a level g is a simple linear translate of the set of equilibrium payoffs from some benchmark level, say, $g = 0$. Consequently, it will be seen that the *set* of emission levels that can arise in equilibrium from level g is identical to those that can arise from equilibrium play at a GHG level of 0. Note the fact that the set of equilibrium possibilities is invariant to the level of g, is perfectly consistent with the possibility that in a particular equilibrium, emission levels vary with g. However, the invariance property will make for a particularly simple characterization of the best and worst equilibria.

Let $\Xi(g)$ denote the *set* of equilibrium payoff vectors with initial state g, i.e., each element of $\Xi(g)$ is the payoff to *some* SPE starting from g.

Theorem 7. *The equilibrium payoff correspondence* Ξ *is linear; there is a compact set* $U \subset \Re^I$ *such that for every initial state* g

$$\Xi(g) = U - \{w_1 g, w_2 g, \ldots, w_I g\},$$

where

$$w_i = \frac{c_i}{1 - \sigma\delta}, \qquad i = 1, \ldots, I.$$

In particular, consider any SPE, any period t *and any history of play up until* t. *Then the payoff vector for the continuation strategies must necessarily be of the form*

$$V - (w_1 g_t, w_2 g_t, \ldots, w_I g_t),$$

where $V \in U$ *(and* g_t *is the state at period* t*).*

The theorem is proved by way of a bootstrap argument. We presume that a (candidate) payoff set has this invariance and show that the linear structure of the model confirms the conjecture. Consequently, we generate another candidate payoff set — which is also state invariant. Then we look for a fixed point of that operator. [In other words, we employ a generalized version of the "Abreu-Pearce-Stachetti operator" to generate the SPE correspondence. We need to generalize the APS argument since that was formulated for repeated games alone; see Dutta and Radner, 2005a.]

4.2 Extreme equilibria

We can now use the result of the previous subsection to characterize the best — and the worst — equilibria in the global climate change game. Consider the *second-best problem* (from initial state g and for a given vector of welfare weights $x = (x_i; i = 1, \ldots, I)$), i.e., the problem of maximizing a weighted sum of *equilibrium payoffs*:

$$\max \sum_{i=1}^{I} x_i V_i(g), \qquad V(g) \in \Xi(g).$$

Note that we consider all possible equilibria, i.e., we consider equilibria that choose to condition on current and past GHG levels as well as equilibria that do not. The result states that the best equilibrium *need not* condition on GHG levels:

Theorem 8. *There exists a constant emission level* $\overline{a} \equiv \overline{a}_1, \overline{a}_2, \ldots, \overline{a}_I$ *— such that no matter what the initial level of GHG, the second-best policy is to emit at the constant rate* $\overline{a}$*. In the event of a deviation from this constant emissions policy by country* i*, play proceeds to* i*'s worst equilibrium. Furthermore, the second-best emission rate is always strictly lower than the BAU rate, i.e.,* $\overline{a} < a^*$*. Above a critical discount factor (less than 1), the second-best rate coincides with the GPO emission rate* $\widehat{a}$*.*

The theorem is attractive for two reasons: first, it says that the best possible equilibrium behavior is no more complicated than BAU behavior; so there is no argument for delaying a treaty (to cut emissions) merely because the status quo is simple. Second, the cut required to implement the second-best policy is an across the board cut — independently of anything else, country i should cut its emissions by the amount $a_i^* - \overline{a}_i$. [Our model operates at the aggregative level alone and, in particular, we do not address the issue of how national governments will implement cuts that they agree to (in the national interest). However, it seems quite likely that an across the board cut will be easier to implement — and will be perceived to be fairer to all — than one which is sensitively tied to levels of GHG.]

Sanctions will be required if countries break with the second-best policy. [A major criticism of the Kyoto accord is that it did not incorporate sanctions and hence would never be carried out. For details, see (Barrett, 2003, Chapter 15), and (Dutta and Radner, 2005a).] Without loss of generality we can restrict attention to the worst such sanction. We turn now to a characterization of this worst equilibrium (for, say, country i). One definition will be useful for this purpose:

Definition 9. *An i-less second-best equilibrium is the solution to a second-best problem in which the welfare weight of i is set equal to zero, i.e., $x_i = 0$.*

By the previous theorem, every such problem has a solution in which on the equilibrium path, emissions are a constant. Denote that emission level $a(x_{-i})$:

Theorem 10. *There exists a "high" emission level $\overline{\overline{a}}(i)$ (with $\sum_{j\neq i} \overline{\overline{a}}_j(i) > \sum_{j\neq i} a_j^*$) and an i-less second-best equilibrium $a(x_{-i})$ such that country i's worst equilibrium is:*

1. Each country emits at rate $\overline{\overline{a}}_j(i)$ for one period (no matter what g is), $j = 1, \ldots, I$.

2. From the second period onwards, each country emits at the constant rate $a_j(x_{-i})$, $j = 1, \ldots, I$.

And if any country k deviates at either stages 1 or 2, play switches to k's worst equilibrium from the very next period after the deviation.

Put another way, for every country i, a sanction is made up of two emission rates, $\overline{\overline{a}}(i)$ and $a(x_{-i})$. The former imposes immediate costs on country i. The way it does so is by increasing the emission levels of countries $j \neq i$. The effect of this is a temporary increase in incremental GHG but due to the irreversibility of gas accumulation, a permament increase in country i's costs, enough of an increase to wipe out any immediate gains that the country might have got from the deviation. Of course this additional emission also increases country j's costs. For the punishing countries, however, this increase is offset by the subsequent permanent change, the switch to the emission vector $a(x_{-i})$, which permamently increases their quota at the expense of country i's.

The fact that there is a temporary loosening of environmental regulations as part of environmental sanctions is reminiscent of GATT rules where tariffs can be temporarily imposed by countries that seek to punish illegitimate trade practices on the part of others.

4.3 Greenhouse trap

In every equilibrium that we have studied so far — BAU, the third and second-best — each country emits at a constant rate regardless of GHG level. Hence, the dynamics of every such equilibrium is also simple; at a constant cumulative rate A, the stock of greenhouse gases (GHGs) converges to a steady-state of $A/(1-\sigma)$. Put yet another way, the current GHG level has no long-term implication.

In this section we demonstrate the richness of the model by identifying some MPE in which current GHG levels matter, MPE in which there is a so-called "greenhouse trap." If the world starts below some critical level of GHG, say $\widetilde{g}$, then it grows no farther than that level. However if the system starts above $\widetilde{g}$ — or somehow crosses into the higher region — then greenhouse gases are trapped into growing — and eventually grow to the BAU steady-state.

The key to these equilibria is a richer interaction between the (Markov) emission levels of country i and the rest of the world. In particular, we will consider emission policies that are Markovian but not constant, say an emission policy (vector) such as $a(g)$. Now country i has an incentive — everything else being equal — to emit in such a fashion that the global stock of GHG grows towards a region where $a_{-i}(.)$ are lower. Of course every country has such an incentive and so each country will wish to drive g_t towards a region where emission levels are low for the other countries. Hence, the conjecture for the group as a whole is that there will be MPE with regions of "abnormally" low emissions and every country will (a) have an incentive to stay in such regions once the system gets there and (b) have an incentive to participate in pushing the system towards such regions.

The conjecture is almost correct. We can show that there are indeed such equilibria — in fact there are many — each of which has such a "good" region of low emissions. Furthermore, from *most* — but not all — initial GHG levels outside this good region, countries will have an incentive to drive the system to the good region. The one additional complication is that in order for there to be a good region there must also be some "bad" region (of high emissions) from which countries do *not* find their way to the good region. The presence of a bad region — and the associated spectre of landing there — is what keeps the countries honest in the good region.

[Note: Readers familiar with repeated games will note the obvious connection with the idea of history-dependent punishments. The point to note though is that these are not history-dependent equilibria since they only depend on the GHG level and not on past emissions. (Hence they are more difficult to construct.) They are also more sparing in terms of informational requirements since they do not require i to condition on (or even know) past emission levels of the other countries.]

To keep the exposition simple, we are only going to present results for the symmetric case. [The first of the two results is easily generalized to the asymmetric case.] Accordingly, when we speak of the Pareto optimal solution we will refer to the symmetric solution — and to avoid clutter we will simply denote that solution by $\widehat{a}$ (with associated steady state $\widehat{g}$). We shall present two results of increasing generality.

Consider the following symmetric Markovian strategy $a(.)$: if the GHG level is below the Pareto optimal steady state $\widehat{g}$, emissions take the game immediately to that state. On the other hand, if the GHG level exceeds $\widehat{g}$, then emissions are at the (high) BAU level of a^*. In other words,

$$a(g) = \frac{\widehat{g} - \sigma g}{I}, \qquad g \leq \widehat{g}$$
$$= a^*, \qquad g > \widehat{g}$$

In the terminology of the immediately preceding discussion, the region below the Pareto optimal steady state $\widehat{g}$ is the "good" region of (relatively) low emission levels whereas the region above is the "bad" region of high emissions. Consider the following condition ("L" for "large"), which says that the BAU emission level is sufficiently larger than the GPO level:

Condition L:
$$\frac{a^*}{\widehat{a}} > \max\left(\frac{I}{I-1}, \frac{1}{1-\sigma}\right),$$

where (recall) I is the number of players and σ is the persistence of CO_2 in the earth's atmosphere.

Remark. Since $I/(I-1)$ approaches 1 for I large, Condition L really boils down to
$$\frac{a^*}{\widehat{a}} \geq \frac{1}{1-\sigma}$$
whenever we have a large number of players. This condition will always hold under standard asymptotic conditions on h. To see this note that simple algebra shows that
$$\frac{h'(\widehat{a})}{h'(a^*)} = I,$$
and hence $a^*/\widehat{a}$ is large whenever I is large.

Theorem 11. *Suppose that Condition L holds. Then there is a cut-off value of the discount factor — say $\widehat{\delta}$ — such that $a(.)$ is an MPE for all $\delta \geq \widehat{\delta}$. In such an equilibrium, the GHG level converges in one period to the Pareto optimal steady state $\widehat{g}$ if the initial level is below $\widehat{g}$ whereas it converges asymptotically to the BAU steady state g^* if the initial level is above $\widehat{g}$.*

The reader might wonder how useful this last theorem is if we suspect that the world is currently already past the Pareto optimal steady state. [It is unclear whether the world is past $\widehat{g}$ or not. Some of the public policy concern is not so much with current levels of GHG as with what level might eventually be attained at the current rates of accumulation.] This leads to the question: are there MPE that have steady states less than the BAU steady state of g^* (but higher than the desired — but unattainable — steady state of $\widehat{g}$)?

We can now demonstrate the existence of MPE that differs in three ways from the one above. First, in addition to the good and bad regions, there will be a "latent good" region — a region of low GHG levels from which the stock will gradually grow till it reaches the good region (whereupon it will stay there). Second, the good region (with emissions lower than the BAU level a^*) will extend beyond the Pareto optimal steady state $\widehat{g}$; in fact it will extend quite close to the BAU steady state g^*. (So no matter where the world is today, however close to the worst possibility, we can still put the brakes on in an incentive-compatible way!) Third, we will considerably weaken Condition L.

Let $\widetilde{g}$ be any GHG level that is higher than $\widehat{g}$ but no more than

$$\frac{(I-1)}{I}g^*;$$

$\widetilde{g}$ will be our candidate low steady state. We will consider Markov strategies in which stocks from a left neighborhood of $\widetilde{g}$ (to be defined shortly) come in one step to $\widetilde{g}$ and stocks above $\widetilde{g}$ exhibit BAU behavior:

$$\begin{aligned}\widetilde{a}(g) &= a^*, & g &> \widetilde{g}\\ &= \frac{\widetilde{g}-\sigma g}{I}, & \widetilde{g}_1 &\leq g \leq \widetilde{g}\end{aligned}$$

where

$$\widetilde{g}_1 \equiv \frac{I-1}{I}\widetilde{g}.$$

For stocks lower than $\widetilde{g}_1$, the "latent good" region, the emission levels will be such that GHG levels grow (gradually) till they get into the $[\widetilde{g}_1, \widetilde{g}]$ region. These emission levels cannot however be solved for in closed form. Instead we can employ a fixed point argument to show the following:

Theorem 12. *Suppose that*

$$\frac{a^*}{\widetilde{a}} > \frac{I}{I-1}.$$

Then there is a cut-off value of the discount factor — say $\widetilde{\delta}$ — such that for all $\delta \geq \widetilde{\delta}$, there is a MPE $\widetilde{a}(.)$ whose behavior above $\widetilde{g}_1$ is as given above. Below $\widetilde{g}_1$, the stock grows although it remains below $\widetilde{g}$, i.e.,

$$\sigma g + I\widetilde{a}(g) \in (g, \widetilde{g}), \qquad \textit{for all } g < \widetilde{g}.$$

In such an equilibrium, the GHG level converges to the steady state $\widetilde{g}$ if the initial level is below it whereas it converges asymptotically to the BAU steady state g^* if the initial level is above $\widetilde{g}$.

5. Concluding remarks

Some extensions of the models and results presented here have been considered in our project, and are reported elsewhere. Others remain for future research. Thus in Dutta and Radner (2005b) we generalize our current model to allow for exogenous population change and demonstrate qualitatively similar theoretical results. A similar approach allows one to allow for exogenous capital accumulation. Sangwon Park (2004) has calibrated several of the theoretical models, and in Dutta and Radner (2004, 2005a) we present some numerical illustrations of the various theoretical results. In Dutta, Park, and Radner (2004) we begin to incorporate endogenous capital accumulation (with technological change). The main question that we hope to address in that model is: (when) does the prevention of global warming slow down the rate of economic growth? A second question that we hope to look at is: (how) does asymmetry in the current level of economic development affect sustainability of agreements about emission cuts? We have some preliminary results on the first question but not a complete solution. Finally, an important piece of unfinished business in our project is the incorporation of uncertainty into our models.

Acknowledgements. A complete list of acknowledgements would be too long for this short summary. However, we must not omit thanking Sangwon Park for his contributions to the project, both theoretical and numerical. We must also thank Geoffrey Heal for his help and encouragement from the early stages of the project onward.

6. Bibliographic note

Background information on the problem of global warming, and references to the relevant literature, can be found in in the references cited in Sections 1 and 5. In particular, as we point out in (Dutta and Radner, 2005a), "A large volume of literature exists on the economics of climate change. A central question there is to determine the level of emissions that is globally optimal. An excellent example of this is Nordhaus and Boyer (2000). A smaller volume of literature emphasizes the need for treaties to be self-enforcing, i.e., the need for a strategic analysis of the problem. (See Barrett (2003) and Finus (2001).) Where we depart from the existing strategic literature is in the dynamic modelling; we allow GHGs to accumulate and stay in the environment for a (possibly long) period of time. (Technically, existing analyses are all static one-shot games or purely repeated games which implies that the state variable, gas stock, remains constant over time.)"

As noted at the end of Section 1, the text of the present paper closely follows the texts of various papers by the authors that are cited in Sections 1 and 5.

In particular, Sections 2 and 3 are based primarily on Dutta and Radner (2004), and Section 4 is based on Dutta and Radner, (2005a).

References

[1] Barrett, S.: Environment and Statecraft: The Strategy of Environmental Treaty-Making. Oxford University Press, Oxford 2003

[2] Finus, M.: Game Theory and International Environmental Cooperation. Edward Elgar, Cheltenham 2001

[3] Dutta, P.K., Radner, R.: Choosing cleaner technologies. (unpublished Working Paper, Economics Department, Columbia University) (2002)

[4] Dutta, P.K., Radner, R.: Self-enforcing climate change treaties. Proceedings of the National Academy of Sciences, USA **101-14**, 5174-5179 (2004)

[5] Dutta, P.K., Radner, R.: A strategic analysis of global warming: Theory and some numbers. (unpublished Working Paper, Economics Department, Stern School, New York University) (2005a)

[6] Dutta, P.K., Radner, R.: Population growth and technological change in a global warming model. Economic Theory, forthcoming (2005b)

[7] Dutta, P.K., Park, S., Radner, R.: Global warming with capital accumulation. in preparation, Columbia University (2004)

[8] Nordhaus, W., Boyer, J.: Warming the World: Economic Models of Global Warming. MIT Press, Cambridge, MA 2000

[9] Park, S.: Calibrating a model of global warming. Columbia University Economics Department (unpublished) (2004)

[10] Radner, R.; A game-theoretic approach to global warming. Woytinsky Lecture, University of Michigan, March 11, 1998 (unpublished Working Paper, Economics Department, Stern School, New York University) (1999)

[11] Raven, P.: Biodiversity and our common future. Bull. of the Amer. Acad. of Arts and Sciences **58**, no.3, 20-25 (2005)